Achievable Performance of Blind Policies in Heavy Traffic

Nikhil Bansal[a,b] n.bansal@tue.nl
Bart Kamphorst[a] b.kamphorst@cwi.nl
Bert Zwart[a,b] b.zwart@cwi.nl

[a] Centrum Wiskunde & Informatica, P.O. Box 94079, 1090 GB Amsterdam, the Netherlands
[b] Technische Universiteit Eindhoven, P.O. Box 513, 5600 MB Eindhoven, the Netherlands

7th January 2016

arXiv:1512.07771v2 [math.PR] 6 Jan 2016

Abstract

For a GI/GI/1 queue, we show that the average sojourn time under the (blind) Randomized Multilevel Feedback algorithm is no worse than that under the Shortest Remaining Processing Time algorithm times a logarithmic function of the system load. Moreover, it is verified that this bound is tight in heavy traffic, up to a constant multiplicative factor. We obtain this result by combining techniques from two disparate areas: competitive analysis and applied probability.

Keywords: GI/GI/1 *queue, average sojourn time, heavy traffic, Shortest Remaining Processing Time, Randomized Multilevel Feedback*

1 Introduction

One of the most relevant and widely studied measures of quality of service in a GI/GI/1 queue is the average sojourn time, also known as response time or flow time, defined as the average time spent by a job from its arrival in the system until its completion [3–5, 7, 9, 16, 20, 22, 29–31]. We consider the most basic setting of a single machine with pre-emption, i.e. jobs can be interrupted arbitrarily and resumed later without any penalty. It is well-known that the Shortest Remaining Processing Time (SRPT) policy, that at any time works on the job with the least remaining processing time, is the optimal policy for every problem instance (or equivalently for every sample path) for minimizing the sojourn time [27]. However, to run SRPT one needs the exact knowledge of all job sizes. This information may not be available in many settings; specifically, jobs sizes may only be known approximately, or may not be known at all [18]. For this reason one may have to be content with more generally applicable policies.

In this paper we are interested in policies that do not require the knowledge of job sizes in their scheduling decisions. We refer to such policies as *blind* policies. More formally, in a blind policy the scheduler is only aware of the existence of a job and how much processing it has received thus far. The size of the job becomes known to the scheduler only when it terminates and leaves the system. Observe that the class of blind policies contains several well-studied policies, such as Processor Sharing (also known as Round Robin) [15], Foreground-Background [21], and First Come First Served [2].

It is natural to ask how much this inability to use the knowledge of sizes can hurt. In particular, how much can the average sojourn time between SRPT and an optimal blind policy differ for a given GI/GI/1 queue? As illustration, let us consider the M/M/1 system. In this system, all blind policies become identical due to the memoryless nature of the job sizes. More precisely, Conway et al. [10] state that any blind policy has average sojourn time $\mathbb{E}[B]/(1-\rho)$, where $\mathbb{E}[B]$ is the average job size and ρ is the load of the system. On the other hand, if job sizes

are known upon arrival, then Bansal [3] derives that under M/M/1/SRPT the average sojourn time $\mathbb{E}[T_{\text{SRPT}}]$ is

$$\mathbb{E}[T_{\text{SRPT}}] = (1 + o(1)) \frac{1}{\log\left(\frac{e}{1-\rho}\right)} \frac{\mathbb{E}[B]}{1-\rho}, \tag{1.1}$$

where $o(1)$ vanishes as ρ approaches one and $\log(\cdot)$ denotes the natural logarithm. That is, the SRPT policy performs better by a factor $\log(e/(1-\rho))$ in heavy traffic. So while SRPT can be arbitrarily better than a blind policy for a M/M/1 system as the load approaches one, this improvement factor grows quite mildly.

The performance of SRPT as a function of load can be dramatically different for more heavy-tailed distributions. Bansal and Gamarnik [4] and Lin et al. [16] show that the growth factor of the average sojourn time in heavy traffic can be much smaller than $1/(1-\rho)$. For example, if the job sizes follow a Pareto(β) distribution with $\beta \in (1, 2)$, then the growth factor of the average sojourn time $\mathbb{E}[T_{\text{SRPT}}]$ is $\mathbb{E}[B] \log(1/(1-\rho))$, up to constant factors depending on β. On the other hand, Kleinrock [15] states that Processor Sharing has average sojourn time $\mathbb{E}[B]/(1-\rho)$. As this example illustrates, it is conceivable that the gap between blind policies and SRPT can be much larger.

Another subfield of computer science where the performance improvement of SRPT over blind policies has been studied is *competitive analysis* [8, 12, 23], that deals with the worst case analysis of algorithms. The study of competitive analysis of blind scheduling policies was initiated by Motwani et al. [19]. They showed that for the problem of minimizing the average sojourn time, no deterministic algorithm can have a better competitive ratio than $\Omega(m^{1/3})$, where m is the number of jobs in an instance. Motwani et al. also showed that no randomized algorithm can have a competitive ratio better than $\Omega(\log(m))$. In a significant breakthrough, Kalyanasundaram and Pruhs [14] gave an elegant and non-trivial randomized algorithm that they called Randomized Multilevel Feedback (RMLF) and proved that it has competitive ratio of at most $O(\log(m) \log(\log(m)))$. Later, Becchetti and Leonardi [6] showed that RMLF is in fact an $O(\log(m))$-competitive randomized algorithm and hence the best possible (up to constant factors). The result in Becchetti and Leonardi [6] is derived under the assumption that job sizes are bounded from below by a strictly positive constant, an assumption which is removed in this paper. Additional background on multilevel algorithms can be found in Kleinrock [15], and an analysis of the mean sojourn time under such algorithms is performed in Aalto and Ayesta [1].

The insights from applied probability and competitive analysis about the relation between blind policies and SRPT can be combined when m is taken as the number of jobs in a regeneration cycle, which has an expected value of the order $1/(1-\rho)$. We make this precise in our main theorem and its proof. The main theorem shows that, for any GI/GI/1 queue, the gap between SRPT and the best blind policy for that system $\mathcal{A}$ is at most $\log(1/(1-\rho))$, up to constant factors. More specifically, we show that this growth factor is a guaranteed upper bound on the gap between SRPT and (an extension of) the RMLF algorithm. That is, we show that

$$\mathbb{E}[T_{\mathcal{A}}] \leq \mathbb{E}[T_{\text{RMLF}}] = O\left(\log\left(\frac{1}{1-\rho}\right) \mathbb{E}[T_{\text{SRPT}}]\right) \tag{1.2}$$

as ρ grows to one. We note that the implementation of the RMLF algorithm does not depend on the distributions of interarrival times and job sizes and is therefore applicable to every GI/GI/1 queue; this property may not hold for an optimal blind policy $\mathcal{A}$ in general.

The second main contribution of this paper is the proof of (1.2) itself. It involves a novel combination of techniques from competitive analysis and applied probability. Using a renewal argument, we consider the average sojourn time $\mathbb{E}[T_{\text{RMLF}}]$ of jobs in a general busy period, and subsequently distinguish two types of busy periods (small and large) by the number of jobs. For small busy periods, we apply a worst-case performance bound of RMLF from the study of

competitive analysis. For large busy periods, we need to derive the heavy traffic behaviour of moments of two functionals: the busy period duration and the number of jobs in a busy period. In particular, we show that the κ-th moment of these functionals behaves like $O((1-\rho)^{1-2\kappa})$ for $\kappa \geq 1$; a result which seems to be new, and may facilitate future instances where competitive analysis and regenerative process theory are combined to obtain information about algorithms under uncertainty. To prove these bounds, we rely on properties of ladder height distributions derived in Asmussen [2] and Lotov [17].

This paper is organized as follows. A detailed model description and notation is introduced in Section 2. Section 3 clarifies the concept of a competitive ratio and describes the RMLF algorithm. Additionally, Section 3.2 relaxes the constraints on RMLF while preserving the competitive ratio. The main result, Theorem 3, is presented in Section 4, whereas the proof is delayed until Section 5. Propositions required for the main theorem are proven in Section 6. Finally, Section 7 concludes the paper.

2 Preliminaries

This section introduces a general framework for sequences of GI/GI/1 queues, so that we may analyse their limiting behaviour in further sections. In particular, the model allows for a heavy traffic analysis of the average sojourn time and various other functionals.

Sequence of queues

Consider a sequence of GI/GI/1 queues, indexed by $n \geq 1$, where jobs arrive one by one with independent and identically distributed (i.i.d.) sizes $B^{(n)}, B_i^{(n)}, i \geq 1$, chosen from a distribution $G_b^{(n)}$. The jobs are then processed by a single server with unit speed. The time between two consecutive job arrivals is given by the i.i.d. interarrival times $A^{(n)}, A_i^{(n)}, i \geq 1$, chosen from a distribution $G_a^{(n)}$. All jobs and interarrival times are positive, i.e. the support of $G_a^{(n)}$ and $G_a^{(n)}$ is contained in $(0, \infty)$.

It is assumed that there exist $\rho_0 \in (0,1)$ and $A_{\min} > 0$ such that $\rho_0 A_{\min} \leq \rho_0 \mathbb{E}[A^{(n)}] \leq \mathbb{E}[B^{(n)}] < \mathbb{E}[A^{(n)}]$ and that both the interarrival times and the job sizes have finite variance for all n. Additionally, it is assumed that $\limsup_{n\to\infty} \mathbb{E}[(A^{(n)})^2] < \infty$ and $\limsup_{n\to\infty} \mathbb{E}[(B^{(n)})^2] < \infty$. The load of the system is denoted by $\rho^{(n)} := \mathbb{E}[B^{(n)}]/\mathbb{E}[A^{(n)}] \in [\rho_0, 1)$, and is interpreted as the fraction of time that the server is busy. As is customary in the literature on heavy traffic analysis [13, 24], we assume $\lim_{n\to\infty} \rho^{(n)} = 1$. The mean amount of work that a server completes between two consecutive arrivals is represented by $\mu^{(n)} := \mathbb{E}[A^{(n)}] - \mathbb{E}[B^{(n)}] = \mathbb{E}[A^{(n)}](1-\rho^{(n)})$. Since a queue can only be initiated by a job arriving in a non-empty system, we pose the additional requirement that for some $\delta > 0$ and $\gamma > 0$ the system satisfies $\mathbb{P}(B^{(n)} - A^{(n)} \geq \delta) \geq \gamma$ for all n.

Example: A key example of a system of queues that fits the above framework is a queue with unit mean job sizes B_i and interarrival times $A_i^{(\rho)} = A_i/\rho, i \geq 1, \rho \in (0,1)$, where the A_i have unit mean. This system experiences a load of $\rho = \mathbb{E}[B_1]/\mathbb{E}[A_1^{(\rho)}]$ and is exposed to heavy traffic as ρ tends to one due to decreasing interarrival times. The described model fits in the framework described above by letting $A_i^{(n)} = A_i/(1-1/n), B_i^{(n)} = B_i$ and $\rho^{(n)} = 1 - 1/n$, and is referred to as the *Example Model*. All further remarks on the Example Model are emphasised by superscripts ρ for all related variables and functionals; e.g. the notation $A_i^{(\rho)}$ will always indicate the i-th interarrival time in this Example Model.

Queueing functionals and netput processes

The sojourn time of a job is the time from the arrival of the job until it completes its service requirement. Given a scheduling policy π, we denote the average sojourn time of a generic job by $\mathbb{E}[T_\pi^{(n)}]$. The steady-state cumulative amount of work in the system is expressed by $V^{(n)}$, which has an atom at zero that corresponds to the times when the server is idle. The steady-state duration of such an idle period is denoted by $I^{(n)}$. Idle periods are ended by the arrival of a new job, which initiates a busy period. A busy period finishes at the earliest subsequent time for which the system is empty again. The steady-state duration of a busy period is represented by $P^{(n)}$, whereas the total number of arrivals during a busy period is symbolized by $N^{(n)}$. Finally, the steady-state cumulative amount of work in the system *at an arrival instance* is represented by $W^{(n)}$. For any non-negative random variable Y, we define a random variable Y_e that is distributed as the excess of Y; i.e. $\mathbb{P}(Y_e \leq x) = \int_0^x \mathbb{P}(Y > x)\,\mathrm{d}x/\mathbb{E}[Y]$.

Next, we define two counting processes in the GI/GI/1 queue under consideration. The first process $N^{(n)}(t) := \inf\{m \in \{1,2,\ldots\} : A_1^{(n)} + \ldots + A_m^{(n)} \geq t\}, t \geq 0$, counts the number of arrivals in t time units, starting from a reference arrival that is also the first count. The second process $\widetilde{N}^{(n)}(t), t \geq 0$ is similar, and only differs by initializing the count at an arbitrary point in time. Specifically,

$$\widetilde{N}^{(n)}(t) = \begin{cases} 0 & \text{if } t < A_e^{(n)}, \\ \inf\{m \in \{1,2,\ldots\} : A_e^{(n)} + A_1^{(n)} + \ldots + A_m^{(n)} \geq t\} & \text{otherwise.} \end{cases}$$

These counting processes allow us to introduce two corresponding netput processes $X^{(n)}(t) = \sum_{i=1}^{N^{(n)}(t)} B_i^{(n)} - t$ and $\widetilde{X}^{(n)}(t) = \sum_{i=1}^{\widetilde{N}^{(n)}(t)} B_i^{(n)} - t$, that quantify the net amount of work that could have been processed by the server in the t time units after an arrival, or respectively after an arbitrary point in time. It is easy to see that $\mathbb{P}(\widetilde{X}^{(n)}(t) > x) \leq \mathbb{P}(X^{(n)}(t) > x)$ for all $t \geq 0$, which will be denoted by $\widetilde{X}^{(n)}(t) \leq_{st} X^{(n)}(t)$ in the remainder of this document.

Similarly, we define two discrete processes that quantify the netput process *at an arrival instance*, denoted by $S_m^{(n)}$ and $\widetilde{S}_m^{(n)}, m \geq 0$. The first process $S_m^{(n)}$ is defined as $S_0^{(n)} := 0, S_m^{(n)} := \sum_{i=1}^m [B_i^{(n)} - A_i^{(n)}]$ and quantifies all the work that the server has received between the arrival of the reference job and the m-th next arrival, minus the work that it could have addressed during this time. The second process starts observing at an arbitrary point in time instead of at the arrival of a reference job. It is defined as $\widetilde{S}_0^{(n)} = -A_e^{(n)}, \widetilde{S}_m^{(n)} = -A_e^{(n)} + \sum_{i=1}^m [B_i^{(n)} - A_i^{(n)}]$. Again, we obtain the relation $\widetilde{S}_m^{(n)} \leq_{st} S_m^{(n)}$. Also, one may verify that $\sup_{t\geq 0} X^{(n)}(t) = \sup_{m=1,2,\ldots} S_m^{(n)}$, and hence by Asmussen [2, Corollary III.6.5] that $\sup_{t\geq 0} X^{(n)}(t) = \sup_{m=0,1,\ldots} S_m^{(n)} \stackrel{d}{=} W^{(n)}$. Here, $\cdot \stackrel{d}{=} \cdot$ denotes equality in distribution. For the remainder of the document, all sums $\sum_{i=1}^0$ are understood to be zero.

Scheduling policies

A scheduling policy π is an algorithm or a rule which specifies how the jobs are executed at any time in the system. For the GI/GI/1 queue under consideration, such a policy prescribes the behaviour of a single server under the relaxation that jobs can be *pre-empted*; that is, jobs can be interrupted at any point during their execution and can be resumed later from this point without any penalty. Of the large class of scheduling policies that apply to this system, we consider those policies π that satisfy the following two criteria (quoted from Wierman and Zwart [31], after Stolyar and Ramanan [28]):

(i) π is *non-anticipative*: a scheduling decision at time t does not depend on information about customers that arrive beyond time t.

(ii) π is *non-learning*: the scheduling decisions cannot depend on information about previous busy periods. That is, a scheduling decision on a sample path cannot change when the history before the current busy period is changed.

Of special interest are those scheduling policies π that additionally obey the following characteristic:

(iii) π is *blind*: the scheduling decisions do not depend on the sizes of the jobs. That is, the scheduling decisions on a sample path up to time t cannot change when the sizes of jobs that have not finished at that time are altered (in such a way that the jobs remain unfinished).

Such policies are very common: First Come First Served (FCFS), Processor Sharing (PS) and Foreground-Background (FB) and all blind policies. On the other hand, policies like Shortest Job First (SJF) or Shortest Remaining Processing Time (SRPT) are *not* blind as they use the knowledge of job sizes while making a scheduling decision. An algorithm satisfying (i) and (iii) is also called *non-clairvoyant*.

We let $\mathcal{A}^{(n)}$ denote a blind policy that minimizes the average sojourn time over the space of all blind policies for the n-th GI/GI/1 queue. In general, $\mathcal{A}^{(n)}$ would depend on the distributions $G_a^{(n)}$ and $G_b^{(n)}$, which specify the GI/GI/1 queue. By RMLF and eRMLF we denote randomized blind algorithms that are formalized in Sections 3.1 and 3.2, respectively. The implementation of the RMLF or eRMLF algorithm does not depend on $G_a^{(n)}$ and $G_b^{(n)}$ and is therefore independent of the system index n.

Asymptotic relations and further functions

We use the standard notation that for functions $f(n)$ and $g(n)$, $f(n) = O(g(n))$ and $f(n) = o(g(n))$ if $\limsup_{n\to\infty} f(n)/g(n) < \infty$ and $\limsup_{n\to\infty} f(n)/g(n) = 0$ respectively. Similarly, $f(n) = \Omega(g(n))$ denotes $\liminf_{n\to\infty} f(n)/g(n) > 0$ and $f(n) = \Theta(g(n))$ is equivalent to $0 < \liminf_{n\to\infty} f(n)/g(n) \le \limsup_{n\to\infty} f(n)/g(n) < \infty$).

Final notational conventions used throughout this document are the floor-function $\lfloor x \rfloor := \sup\{m \in \mathbb{N} : m \le x\}$ and the indicator function $\mathbb{1}\{\textit{logical expression}\}$ that takes value 1 if the logical expression is true, and value 0 otherwise. The $\log(\cdot)$-function denotes the natural logarithm.

3 Competitive analysis of scheduling policies

In this section, we describe some relevant definitions and results from the area of competitive analysis, which deals with the worst case analysis of algorithms. We restrict our presentation here only to the competitive analysis of scheduling algorithms, and we introduce both the original RMLF algorithm and an extension of this algorithm for later purposes.

A scheduling problem instance $\mathcal{I}$ consists of a collection of jobs specified by their sizes and their arrival times. We say that an instance has size m, i.e. $|\mathcal{I}| = m$, if it consists of m jobs. For an instance $\mathcal{I}$, we denote the optimal average sojourn time possible for this instance by $T_{opt}(\mathcal{I})$, which for our purposes is same as $T_{\mathrm{SRPT}}(\mathcal{I})$.

For a deterministic algorithm π, we let $T_\pi(\mathcal{I})$ denote the average sojourn time when the instance $\mathcal{I}$ is executed according to the algorithm π. We say that the algorithm π has competitive ratio $c(m)$ if

$$\sup_{\mathcal{I}:|\mathcal{I}|\le m} \frac{T_\pi(\mathcal{I})}{T_{opt}(\mathcal{I})} \le c(m).$$

Thus, the competitive ratio of an algorithm (possibly a function of m), is the worst case ratio over all input instances of length at most m of the sojourn time achieved by π and the

optimal sojourn time on that instance. Observe that the definition of the competitive ratio is rather strict, in that even if a algorithm is close to optimal on all but one input instance, its competitive ratio will be lower bounded by its performance on the bad input instance.

For this purpose it is useful to consider randomized algorithms. A randomized algorithm $\widetilde{\pi}$ can toss coins internally and base its decisions on this outcome of these internal random variables. Such an algorithm can thus be viewed as a probability distribution over deterministic algorithms π_i [8]. It then follows that the average sojourn time of instance $\mathcal{I}$ under a randomized algorithm $\widetilde{\pi}$ equals $T_{\widetilde{\pi}}(\mathcal{I}) = \mathbb{E}_i[T_{\pi_i}(\mathcal{I})]$, where the expectation is over the internal random choices of the algorithm. We say that $\widetilde{\pi}$ has competitive ratio $c(m)$ if

$$\sup_{\mathcal{I}:|\mathcal{I}|\leq m} \frac{T_{\widetilde{\pi}}(\mathcal{I})}{T_{opt}(\mathcal{I})} = \sup_{\mathcal{I}:|\mathcal{I}|\leq m} \frac{\mathbb{E}_i[T_{\pi_i}(\mathcal{I})]}{T_{opt}(\mathcal{I})} \leq c(m). \tag{3.1}$$

Observe that the expectation is only over the random choices made by the algorithm, and the competitive ratio is still determined by the worst possible instance. However, the competitive ratio of a randomized algorithm can be substantially lower, e.g. in situations where no single (deterministic) algorithm is good for all instances, but a suitable combination of algorithms is close to optimal for all instances.

3.1 Randomized Multilevel Feedback algorithm

This section introduces Kalyanasundaram and Pruhs [14]'s Randomized Multilevel Feedback (RMLF) algorithm. As the name suggests, it is a randomized version of the Multilevel Feedback (MLF) algorithm by Corbató et al. [11]. Both algorithms are blind and can therefore only learn the size of a job upon completion.

The general idea of both MLF and RMLF is to prioritize potential short jobs (e.g. jobs that have not received much service) and reduce the priority of a job as it receives more service. This prioritisation is embodied by assigning every job J_j to a virtual high priority queue Q_i, and move it to a lower priority queue Q_{i+1} once it has received $T_{i,j}$ units of service. The performance of the algorithm may suffer from a poor choice of the so-called targets $T_{i,j}$; in particular, if the job sizes are slightly above the targets, then jobs are moved to lower priority queues just prior to completion. The improvement of RMLF over MLF is due to randomization of the targets, thereby reducing the possibility of such events.

We now provide a mathematical representation of the RMLF algorithm. Assume first that there is a universal lower bound of on the job sizes in every instance $\mathcal{I}$, say with value 2. For every instance of size m, the j-th job J_j is released at time r_j and has size B_j. The process $w_j(t)$ denotes the amount of time that RMLF has run J_j before time t. For some symbolic constant θ, fixed at $\theta := 12$, we define the exponentially distributed variables β_j with $\mathbb{P}(\beta_j \leq x) = 1 - \exp[-\theta x \ln j]$. Finally, the targets are defined as $T_{i,j} = 2^i \max\{1, 2 - \beta_j\}$ for all $i \geq 1, 1 \leq j \leq m$. RMLF is then formalized in Figure 1.

Kalyanasundaram and Pruhs [14] proved that the RMLF algorithm is $O(\log(m)\log(\log(m)))$-competitive. This result was later strengthened by Becchetti and Leonardi [7] to a competitive ratio of $O(\log(m))$:

Theorem 1. *The* RMLF *algorithm is* $\log(m)$*-competitive. That is,*

$$\mathbb{E}[T_{\mathrm{RMLF}}(\mathcal{I})] \leq C_1 \log(m) T_{\mathrm{SRPT}}(\mathcal{I}) \tag{3.2}$$

for all instances $\mathcal{I}$ *of size at most* m *and a universal constant* C_1*.*

The competitive ratio lower bound of $\Omega(\log(m))$ as proclaimed by Motwani et al. [19] implies that, up to multiplicative factors, this is the best bound possible for randomized algorithms in

Algorithm RMLF**:** At all times the collection of released, but uncompleted, jobs are partitioned into queues, $Q_0, Q_1, \ldots$ We say that Q_i is lower than Q_j for $i < j$. [For] each job $J_j \in Q_i, T_{i,j} \in [2^i, 2^{i+1}]$ when it entered Q_i. RMLF maintains the invariant that it is always running the job at the front of the lowest non-empty queue.
When a job J_h is released at time r_h, RMLF takes the following actions:

- Job J_h is enqueued on Q_0.
- The target $T_{0,h}$ is set to $\max\{1, 2 - \beta_j\}$.
- If, just prior to r_h, it was the case that Q_0 was empty, and that RMLF was running a job J_j, RMLF then takes the following actions:
  - Job J_j is pre-empted. Note that J_j remains at the front of its queue.
  - RMLF begins running J_h.

If at some time t, a job $J_j \in Q_{i-1}$ is being run when $w_j(t)$ becomes equal to $T_{i-1,j}$, then RMLF takes the following actions:

- Job J_j is dequeued from Q_{i-1}.
- Job J_j is enqueued on Q_i.
- The target $T_{i,j}$ is set to $2T_{i-1,j} = 2^i \max\{1, 2 - \beta_j\}$.

Whenever a job is completed, it is removed from its queue. Recall that RMLF maintains the invariant that it is always running the job at the front of the lowest non-empty queue.

Figure 1: Formal statement of eRMLF algorithm.

the current model. Note that this competitive ratio is significantly lower than the best possible ratio for deterministic algorithms: $\Omega(m^{1/3})$.

In the next section we propose a variant on RMLF that makes the assumption of a universal lower bound on job sizes obsolete.

3.2 Extending the RMLF algorithm

In a general GI/GI/1 queue there may not be a strictly positive lower bound on the job sizes. The RMLF algorithm is not directly applicable in that case. This problem is solved in an extension of the RMLF algorithm, which we will refer to as the eRMLF algorithm. The eRMLF algorithm defines queues $\widetilde{Q}_1, \widetilde{Q}_2, \ldots$ that are identical to the queues $Q_1, Q_2, \ldots$ of the RMLF algorithm, but splits the first queue Q_0 into many queues $\widetilde{Q}_0, \widetilde{Q}_{-1}, \ldots$ Additionally, it considers a "new job" queue $\widetilde{Q}^*$. The concept of the eRMLF algorithm is described below; the formal statement is presented in Appendix A.

Let a problem instance $\widetilde{\mathcal{I}}$ for eRMLF be given. A target $\widetilde{T}_{*,j} = 2^{z^*} \max\{1, 2 - \widetilde{\beta}_j\}$ is assigned to every job $\widetilde{J}_j$ upon arrival, where $\widetilde{\beta}_j$ is an exponentially distributed random variable and $z^* \in \mathbb{Z}$ depends on the current state of the system. When the target has been assigned to the new job, it receives service in $\widetilde{Q}^*$ until the obtained service equals the target, or until a new job arrives. Once either of these events occurs, the job in $\widetilde{Q}^*$ is assigned to a queue $\widetilde{Q}_z, z \in \mathbb{Z}$.

If there are no jobs in queue $\widetilde{Q}^*$, the eRMLF algorithm serves the queues $\widetilde{Q}_z$ in a similar fashion as the RMLF algorithm. Moreover, at any time the problem instance $\widetilde{\mathcal{I}}$ can be converted to a problem instance $\mathcal{I}$ for RMLF by a scaling argument, and under this scaling the sojourn times of all jobs are identical for both algorithms. From this perspective, it is only natural that eRMLF inherits the competitive ratio of RMLF:

Theorem 2. *The* eRMLF *algorithm is* $\log(m)$*-competitive. That is,*

$$\mathbb{E}[T_{\text{eRMLF}}(\mathcal{I})] \leq C_1 \log(m) T_{\text{SRPT}}(\mathcal{I}) \tag{3.3}$$

for all instances $\mathcal{I}$ *of size at most* m *for a universal constant* C_1*. This constant is identical to the constant* C_1 *in Theorem 1.*

The proof of Theorem 2 is given in Appendix A.

4 Main result

We are now ready to present the main result, Theorem 3. The main result states that, for any GI/GI/1 queue, the average sojourn time under SRPT is at most a factor $\log(1/(1-\rho^{(n)}))$ better than that under eRMLF:

Theorem 3. *For any* GI/GI/1 *queue, the* eRMLF *algorithm satisfies the relation*

$$\mathbb{E}[T^{(n)}_{\text{eRMLF}}] = O\left(\log\left(\frac{1}{1-\rho^{(n)}}\right)\mathbb{E}[T^{(n)}_{\text{SRPT}}]\right) \tag{4.1}$$

as $n \to \infty$*, provided that* $\sup_{n=1,2,\ldots} \mathbb{E}[(B^{(n)})^{\alpha}] < \infty$ *for some* $\alpha > 2$*.*

Proof. The proof of the theorem is postponed until the next section. It relies on techniques from both competitive analysis and applied probability. □

As a consequence, the blind policy $\mathcal{A}^{(n)}$ that minimizes the average sojourn time for the n-th system also satisfies the above performance bound. We emphasise the fact that the implementation of eRMLF does not depend on the interarrival and job size distributions, whereas this may not be true for the optimal blind policy $\mathcal{A}^{(n)}$. This property may pose a considerable advantage over a system-dependent optimal blind policy with similar mean performance, e.g. if the input distributions are only approximately known. The section is concluded with some remarks:

Remark 1. Recall that the average sojourn time under any blind policy in an M/M/1 queue is $\mathbb{E}[B^{(n)}]/(1-\rho^{(n)})$, whereas the average sojourn time under SRPT is

$$\mathbb{E}[T^{(n)}_{\text{SRPT}}] = (1+o(1))\frac{1}{\log\left(\frac{e}{1-\rho^{(n)}}\right)}\frac{\mathbb{E}[B^{(n)}]}{1-\rho^{(n)}}, \tag{4.2}$$

where $o(1)$ vanishes as n tends to infinity [5]. In particular, our result is tight up to a multiplicative factor.

Remark 2. As noted in the introduction, the growth rate of $\mathbb{E}[T^{(n)}_{\text{SRPT}}]$ scales like $\log(1/(1-\rho^{(n)}))$ if the job sizes follow a Pareto(β) distribution with $\beta < 2$ (e.g. $\mathbb{P}(B > x) = x^{-\beta}, x \geq 1$) [16]. Since the Pareto distribution has a Decreasing Failure Rate (DFR), Nuyens and Wierman [21] and Righter et al. [26] claim that the optimal blind policy $\mathcal{A}^{(n)}$ for this system is FB, for which the average sojourn time grows with the same rate [4, 21]. Likewise, FB has the same growth rate as SRPT for Pareto(β) distributions if $\beta > 2$ (up to multiplicative factors). This shows that, for specific distributions $G^{(n)}_a$ and $G^{(n)}_b$, policies $\mathcal{A}^{(n)}$ may be found which outperform eRMLF. However, these policies may only perform well for certain classes of input distributions (e.g. FB if $G^{(n)}_b$ has DFR) or they may depend on complete knowledge of the distributions $G^{(n)}_a$ and $G^{(n)}_b$.

5 Proof of the main theorem

The current section presents the proof of Theorem 3. We present an overview of the proof before descending into the technicalities.

5.1 Strategy

The key idea of the proof is that the competitive ratio of the eRMLF algorithm provides an upper bound on the performance of eRMLF. Specifically, it guarantees an upper bound of $O(\log(m))$ on the ratio of the performance of eRMLF over the performance of SRPT, for instances of length at most m. Unfortunately, a general GI/GI/1 queue corresponds to an infinite problem instance and hence the competitive ratio result can not be applied directly. Using the fact that a GI/GI/1 queue is a renewal process, it suffices to analyse individual busy periods rather than the infinite problem instance. Even so, regarding every busy period as an individual problem instance does not bound the problem instance length.

One way to circumvent the unbounded problem instances is by discriminating between "small" busy periods with at most $N_0^{(n)}$ jobs, and "large" busy periods. Busy periods with at most $N_0^{(n)}$ jobs can be analysed with the competitive ratio, yielding a bound of $O(\log(N_0^{(n)}))$. This leaves us with the analysis of large busy periods.

Since the GI/GI/1 queue induces a distribution over problem instances, the probability of experiencing busy periods with at least $N_0^{(n)}$ jobs can be made arbitrarily small by choosing the threshold $N_0^{(n)}$ properly. The combined sojourn time of all jobs in such a large busy period is dominated by the product of the number of jobs $N^{(n)}$ in the busy period and the duration $P^{(n)}$ of the busy period. Therefore, the contribution of large busy periods to the overall average sojourn time is at most $\mathbb{E}[N^{(n)}P^{(n)}\mathbb{1}\{N^{(n)} > N_0^{(n)}\}]/\mathbb{E}[N^{(n)}]$. We will show that, for an appropriate choice of $N_0^{(n)}$, the contribution of the large busy periods is $o(\log(N_0^{(n)}))$.

The second part of this section formalizes the above strategy. In the analysis of $\mathbb{E}[N^{(n)}P^{(n)}\mathbb{1}\{N^{(n)} > N_0^{(n)}\}]$ we greatly rely on Hölder's inequality for decoupling the given expectation into individual moments of $P^{(n)}$ and $N^{(n)}$. The behaviour of these moments is then subject of Propositions 1 and 2, both of which are proven in Section 6.

5.2 Small and large busy periods

We will begin by specifying the threshold that distinguishes small and large busy periods based on the number of jobs. Fix $r \in \left(\frac{\alpha}{\alpha-1}, 2\right)$ and $f > \frac{4+2r}{2-r}$. The threshold $N_0^{(n)}$ is now defined as $N_0^{(n)} := (1-\rho^{(n)})^{-f}$.

Let $T^{(n)}_{\text{eRMLF},i}, T^{(n)}_{\text{SRPT},i}, i = 1, \ldots, N^{(n)}$, be the sojourn time of job i under algorithm eRMLF and SRPT, respectively. Using the fact that a GI/GI/1 queue is a renewal process, we only need to consider a general busy period when analysing the average sojourn time:

$$\mathbb{E}[T^{(n)}_{\text{eRMLF}}] = \frac{1}{\mathbb{E}[N^{(n)}]}\mathbb{E}\left[\sum_{i=1}^{N^{(n)}} T^{(n)}_{\text{eRMLF},i}\right]. \tag{5.1}$$

Discriminating between small and large busy periods then yields

$$\mathbb{E}[T^{(n)}_{\text{eRMLF}}] = \frac{1}{\mathbb{E}[N^{(n)}]}\mathbb{E}\left[\sum_{i=1}^{N^{(n)}} T^{(n)}_{\text{eRMLF},i}\mathbb{1}\{N^{(n)} \le N_0^{(n)}\}\right] + \frac{1}{\mathbb{E}[N^{(n)}]}\mathbb{E}\left[\sum_{i=1}^{N^{(n)}} T^{(n)}_{\text{eRMLF},i}\mathbb{1}\{N^{(n)} > N_0^{(n)}\}\right]. \tag{5.2}$$

As described in the strategy, we will bound the first term by means of the competitive ratio of eRMLF and the second term by Hölder's inequality. These analyses are subject of the following two subsections.

5.3 Small busy periods: competitive ratio

The first term in (5.2) considers busy periods with at most $N_0^{(n)}$ jobs. Theorem 2 ensures that, for any problem instance $\mathcal{I}$ with $N^{(n)} \leq N_0^{(n)}$ jobs, the mean sojourn time $T_{\text{eRMLF}}(\mathcal{I})$ is bounded by $C_1 \log(N_0^{(n)}) T_{\text{SRPT}}(\mathcal{I})$. In particular

$$\begin{aligned}\frac{1}{\mathbb{E}[N^{(n)}]}\mathbb{E}\left[\sum_{i=1}^{N^{(n)}} T^{(n)}_{\text{eRMLF},i}\mathbb{1}\{N^{(n)} \leq N_0^{(n)}\}\right] &\leq \frac{C_1}{\mathbb{E}[N^{(n)}]}\log(N_0^{(n)})\mathbb{E}\left[\sum_{i=1}^{N^{(n)}} T^{(n)}_{\text{SRPT},i}\mathbb{1}\{N^{(n)} \leq N_0^{(n)}\}\right] \\ &\leq \frac{C_1}{\mathbb{E}[N^{(n)}]}\log(N_0^{(n)})\mathbb{E}\left[\sum_{i=1}^{N^{(n)}} T^{(n)}_{\text{SRPT},i}\right] \\ &= C_1 \log(N_0^{(n)})\mathbb{E}[T^{(n)}_{\text{SRPT}}].\end{aligned}$$

The proof is complete once we show that the second term in (5.2) is dominated by $\log(N_0^{(n)})\mathbb{E}[T^{(n)}_{\text{SRPT}}]$ as $n \to \infty$.

5.4 Large busy periods: Hölders inequality

A bound on the sojourn time of an individual job is provided by the duration $P^{(n)}$ of the busy period. Therefore, the second term in (5.2) is bounded by

$$\frac{1}{\mathbb{E}[N^{(n)}]}\mathbb{E}\left[\sum_{i=1}^{N^{(n)}} T^{(n)}_{\text{eRMLF},i}\mathbb{1}\{N^{(n)} > N_0^{(n)}\}\right] \leq \frac{1}{\mathbb{E}[N^{(n)}]}\mathbb{E}\left[N^{(n)}P^{(n)}\mathbb{1}\{N^{(n)} > N_0^{(n)}\}\right]. \tag{5.3}$$

We will temporarily discard the factor $1/\mathbb{E}[N^{(n)}]$. The functionals $N^{(n)}$ and $P^{(n)}$ are dependent, which makes an exact analysis of the expectation troublesome. This complication is avoided by applying Hölder's inequality, which allows us to approximate the dependent expectation by two independent expectations. In particular, for $s := \frac{r}{r-1} \in (2, \alpha)$ we have $\frac{1}{r} + \frac{1}{s} = 1$ and hence

$$\mathbb{E}\left[\sum_{i=1}^{N^{(n)}} T^{(n)}_{\text{eRMLF},i}\mathbb{1}\{N^{(n)} > N_0^{(n)}\}\right] \leq \mathbb{E}[(P^{(n)})^{\frac{r}{r-1}}]^{\frac{r-1}{r}}\mathbb{E}[(N^{(n)})^r\mathbb{1}\{N^{(n)} > N_0^{(n)}\}]^{\frac{1}{r}}. \tag{5.4}$$

Applying Hölder's inequality once more with parameters $\frac{2}{r}$ and $\frac{2}{2-r}$, we get

$$\mathbb{E}\left[\sum_{i=1}^{N^{(n)}} T^{(n)}_{\text{eRMLF},i}\mathbb{1}\{N^{(n)} > N_0^{(n)}\}\right] \leq \mathbb{E}[(P^{(n)})^{\frac{r}{r-1}}]^{\frac{r-1}{r}}\mathbb{E}[(N^{(n)})^2]^{\frac{1}{2}}\mathbb{P}(N^{(n)} > N_0^{(n)})^{\frac{2-r}{2r}}. \tag{5.5}$$

Finally, the tail probability of $N^{(n)}$ is bounded by Markov's inequality. Returning the factor $1/\mathbb{E}[N^{(n)}]$ gives the following upper bound for the second term in (5.2):

$$\frac{1}{\mathbb{E}[N^{(n)}]}\mathbb{E}\left[\sum_{i=1}^{N^{(n)}} T^{(n)}_{\text{eRMLF},i}\mathbb{1}\{N^{(n)} > N_0^{(n)}\}\right] \leq \mathbb{E}[(P^{(n)})^{\frac{r}{r-1}}]^{\frac{r-1}{r}}\mathbb{E}[(N^{(n)})^2]^{\frac{1}{2}}\frac{\mathbb{E}[N^{(n)}]^{\frac{2-r}{2r}-1}}{(N_0^{(n)})^{\frac{2-r}{2r}}}. \tag{5.6}$$

The analysis of the average sojourn time for large busy periods has now been reduced to the analysis of moments of $N^{(n)}$ and $P^{(n)}$. The following two propositions quantify the behaviour of these moments.

Proposition 1. *Assume* $\sup_{n=1,2,\ldots} \mathbb{E}[(B^{(n)})^\alpha] < \infty$ *for some* $\alpha \geq 2$. *Then*

$$\mathbb{E}[(P^{(n)})^\kappa] = O\left((1-\rho^{(n)})^{-(2\kappa-1)}\right). \tag{5.7}$$

for all $\kappa \in [1, \alpha]$. *Moreover,* $\mathbb{E}[P^{(n)}] = \Theta\left((1-\rho^{(n)})^{-1}\right)$.

Proposition 2. *Assume* $\sup_{n=1,2,\ldots} \mathbb{E}[(B^{(n)})^\alpha] < \infty$ *for some* $\alpha \geq 2$. *Then*

$$\mathbb{E}[(N^{(n)})^\kappa] = O\left((1-\rho^{(n)})^{-(2\kappa-1)}\right). \tag{5.8}$$

for all $\kappa \in [1, \alpha]$. *Moreover,* $\mathbb{E}[N^{(n)}] = \Theta\left((1-\rho^{(n)})^{-1}\right)$.

Proof. Both propositions are proven in Section 6. □

Remark 3. When applied to the Example Model, the above propositions state that both $\mathbb{E}[(P^{(\rho)})^\kappa]$ and $\mathbb{E}[(N^{(\rho)})^\kappa]$ are bounded from above by $C_2(1-\rho)^{-(2\kappa-1)}$ uniformly in ρ, for some constant C_2. The integer moments of the busy period duration in an M/GI/1 queue can alternatively be calculated directly from its Laplace-Stieltjes transform, yielding $\mathbb{E}[P^{(\rho)}] = \frac{\mathbb{E}[B]}{1-\rho}$ and $\mathbb{E}[(P^{(\rho)})^2] = \frac{\mathbb{E}[B^2]}{(1-\rho)^3}$. One may therefore conclude that the asymptotic behaviour of the bound in Proposition 1 is in fact sharp for the first two moments of the busy period duration $P^{(\rho)}$ in the Example Model.

From Propositions 1 and 2 it follows that, for some constant C_3,

$$\begin{aligned}
&\frac{1}{\mathbb{E}[N^{(n)}]}\mathbb{E}\left[P^{(n)}N^{(n)}\mathbb{1}\{N^{(n)} > N_0^{(n)}\}\right] \\
&\qquad \leq C_3(1-\rho^{(n)})^{-\left(2\frac{r}{r-1}-1\right)\frac{r-1}{r}}(1-\rho^{(n)})^{-\frac{3}{2}}(1-\rho^{(n)})^{-\frac{2-r}{2r}+1}(1-\rho^{(n)})^{\frac{2-r}{2r}f} \\
&\qquad = C_3(1-\rho^{(n)})^{-1-\frac{1}{r}}(1-\rho^{(n)})^{-\frac{3}{2}}(1-\rho^{(n)})^{-\frac{1}{r}+\frac{3}{2}}(1-\rho^{(n)})^{\frac{2-r}{2r}f} \\
&\qquad = C_3(1-\rho^{(n)})^{\frac{2-r}{2r}f-\frac{2+r}{r}}.
\end{aligned} \tag{5.9}$$

Comparing this to $\log(N_0)\mathbb{E}[T^{(n)}_{\mathrm{SRPT}}]$ and noting that $\mathbb{E}[T^{(n)}_{\mathrm{SRPT}}] \geq \mathbb{E}[B^{(n)}]$ yields

$$\begin{aligned}
\limsup_{n\to\infty} \frac{\frac{1}{\mathbb{E}[N^{(n)}]}\mathbb{E}\left[P^{(n)}N^{(n)}\mathbb{1}\{N^{(n)} > N_0^{(n)}\}\right]}{\log(N_0^{(n)})\mathbb{E}[T^{(n)}_{\mathrm{SRPT}}]} &\leq \limsup_{n\to\infty} \frac{C_3(1-\rho^{(n)})^{\frac{2-r}{2r}f-\frac{2+r}{r}}}{a\log\left(\frac{1}{1-\rho^{(n)}}\right)\mathbb{E}[B^{(n)}]} \\
&\leq \limsup_{n\to\infty} \frac{C_3(1-\rho^{(n)})^{\frac{2-r}{2r}f-\frac{2+r}{r}}}{a\log\left(\frac{1}{1-\rho^{(n)}}\right)\rho_0 A_{\min}},
\end{aligned}$$

which tends to zero as $n \to \infty$. This completes the proof of Theorem 3.

6 Moments of busy period functionals

The moments of the busy period length $P^{(n)}$ and jobs in a busy period $N^{(n)}$ were bounded in Propositions 1 and 2. Some of the required machinery is summarized in Lemmas 1 and 2, which concern the first moment of $N^{(n)}$ and moments of the work $W^{(n)}$ in the system at an arrival, respectively. Recall that $f(n) = \Theta(g(n))$ if $0 < \liminf_{n\to\infty} f(n)/g(n) \leq \limsup_{n\to\infty} f(n)/g(n) < \infty$.

Lemma 1. *The following two relations hold for the steady-state number of customers during a busy period* $N^{(n)}$ *and the steady-state idle time duration* $I^{(n)}$ *as* $n \to \infty$*:*

$$(1-\rho^{(n)})\mathbb{E}[N^{(n)}] = \Theta(1) \text{ and } \mathbb{E}[I^{(n)}] = \Theta(1). \tag{6.1}$$

Moreover, both relations in (6.1) *hold true if the finiteness condition* $\sup_{n=1,2,\ldots} \mathbb{E}[(B^{(n)})^2] < \infty$ *is relaxed to* $\sup_{n=1,2,\ldots} \mathbb{E}[B^{(n)}] < \infty$.

Proof. We will prove the relation $\mu^{(n)}\mathbb{E}[N^{(n)}] = \Theta(1)$, noting that $\mu^{(n)} = \mathbb{E}[A^{(n)}](1-\rho^{(n)}) = \Theta(1-\rho^{(n)})$. Proposition X.3.1 in Asmussen [2], stating

$$\mathbb{E}[I^{(n)}] = \mu^{(n)}\mathbb{E}[N^{(n)}], \tag{6.2}$$

then implies that this is equivalent to the relation $\mathbb{E}[I^{(n)}] = \Theta(1)$.

The upper bound follows directly from $\sup_{n\in\{1,2,\ldots\}} \mathbb{E}[(A^{(n)})^2] < \infty$ and Theorem 2 in Lotov [17], which claims

$$\mu^{(n)}\mathbb{E}[N^{(n)}] \leq C_4 \tag{6.3}$$

for some constant C_4 and all n.

Our model properties state that there exist constants $\delta > 0, \gamma > 0$ such that $\mathbb{P}(B^{(n)} - A^{(n)} \geq \delta) \geq \gamma$ for all n. From inequality (2) in Lotov [17] it then follows that

$$\mu^{(n)}\mathbb{E}[N^{(n)}] \geq \int_0^\infty x \,\mathrm{d}\mathbb{P}(B^{(n)} - A^{(n)} \leq x) \geq \delta\gamma$$

for all n. This completes the proof.

□

Lemma 2. *Let $p > 0$ and define $q = \max\{2, p+1\}$. Assume that $\sup_{n=1,2,\ldots} \mathbb{E}[(B^{(n)})^q] < \infty$. Then*

$$\limsup_{n\to\infty} (1-\rho^{(n)})^p \mathbb{E}[(W^{(n)})^p] < \infty. \tag{6.4}$$

Remark 4. Consider the Example Model, and assume that jobs are served according to the FIFO discipline. Then $W^{(\rho)}$ is just the waiting time of a job, and hence for some constant C_5 the average sojourn time $\mathbb{E}[T^{(\rho)}_{\mathrm{FIFO}}] = \mathbb{E}[W^{(\rho)}_{\mathrm{FIFO}}] + \mathbb{E}[B] \leq \frac{C_5}{1-\rho} + \mathbb{E}[B]$ scales no worse than $1/(1-\rho)$. Lemma 2 provides bounds on the work $W^{(\rho)}$ at an arrival for more general moments, provided that a sufficiently high moment of the job size distribution exists.

Proof. Since $\sup_{n=1,2,\ldots} \mathbb{E}[A^{(n)}] < \infty$, relation (6.4) is equivalent to

$$(\mu^{(n)})^p \mathbb{E}[(W^{(n)})^p] < \infty. \tag{6.5}$$

It is this relation that is proven below.

Assume $p \geq 1$ and let $E^{(n)}, E^{(n)}_i, i \geq 1$, be independent exponentially distributed random variables with mean

$$\mathbb{E}[E^{(n)}] = \frac{\mathbb{E}[A^{(n)}] + \mathbb{E}[B^{(n)}]}{2} \leq \mathbb{E}[A^{(n)}],$$

noting that $\sup_{n=1,2,\ldots} \mathbb{E}[E^{(n)}] \leq \sup_{n=1,2,\ldots} \mathbb{E}[A^{(n)}] < \infty$.

By Asmussen [2, Corollary III.6.5] and subadditivity of suprema, $W^{(n)}$ is upper bounded as

$$\begin{aligned} W^{(n)} &\stackrel{d}{=} \sup_{m\in\{0,1,2,\ldots\}} \sum_{i=1}^{m} \left[B_i^{(n)} - A_i^{(n)}\right] \\ &\leq \sup_{m\in\{0,1,2,\ldots\}} \sum_{i=1}^{m} \left[B_i^{(n)} - E_i^{(n)}\right] + \sup_{m\geq 0} \sum_{i=1}^{m} \left[E_i^{(n)} - A_i^{(n)}\right] \\ &=: W_1^{(n)} + W_2^{(n)}, \end{aligned}$$

where $W_1^{(n)}$ can be interpreted as the total work in an M/GI/1 queue as observed by an arrival, and $W_2^{(n)}$ as the total work in an GI/M/1 queue as observed by an arrival. As a consequence,

$\mathbb{P}(W^{(n)} > x) \le \mathbb{P}(W_1^{(n)} + W_2^{(n)} > x) \le \mathbb{P}(W_1^{(n)} > x/2) + \mathbb{P}(W_2^{(n)} > x/2)$ and thus

$$\begin{aligned}\mathbb{E}[(W^{(n)})^p] &= p\int_0^\infty x^{p-1}\mathbb{P}(W^{(n)} > x)\,\mathrm{d}x \\ &\le p\int_0^\infty x^{p-1}\mathbb{P}(W_1^{(n)} > x/2)\,\mathrm{d}x + p\int_0^\infty x^{p-1}\mathbb{P}(W_2^{(n)} > x/2)\,\mathrm{d}x \\ &= 2^p\left(\mathbb{E}[(W_1^{(n)})^p] + \mathbb{E}[(W_2^{(n)})^p]\right).\end{aligned}$$

First, we consider $W_1^{(n)}$. Define the geometrically distributed random variable $K_1^{(n)}$ with support $\{0, 1, \ldots\}$ and fail parameter

$$\xi_1^{(n)} := \frac{\mathbb{E}[B^{(n)}]}{\mathbb{E}[E^{(n)}]} = \frac{2\mathbb{E}[B^{(n)}]}{\mathbb{E}[A^{(n)}] + \mathbb{E}[B^{(n)}]},$$

and let $B_e^{(n)}, B_{e,i}^{(n)}, i \ge 1$ be i.i.d. random variables that are distributed as the excess job size: $\mathbb{P}(B_{e,i}^{(n)} \le x) = \int_0^x \mathbb{P}(B^{(n)} > t)\,\mathrm{d}t/\mathbb{E}[B^{(n)}]$. Theorem VIII.5.7 in Asmussen [2] presents a random sum representation of the functional $W_1^{(n)}$ in terms of $K_1^{(n)}$ and $B_{e,i}^{(n)}$:

$$W_1^{(n)} \stackrel{d}{=} \sum_{i=1}^{K_1^{(n)}} B_{e,i}^{(n)}.$$

Since $f(x) = x^p$ is a convex function for all $p \ge 1$, Lemma 5 in Remerova et al. [25] implies

$$\mathbb{E}[(W_1^{(n)})^p] \le \mathbb{E}[(K_1^{(n)})^p]\mathbb{E}[(B_e^{(n)})^p]. \tag{6.6}$$

The conditions of the lemma ensure that the p-th moment of $B_e^{(n)}$ is finite as $n \to \infty$:

$$\begin{aligned}\mathbb{E}[(B_e^{(n)})^p] &= \int_0^\infty x^p\,\mathrm{d}\mathbb{P}(B_e^{(n)} \le x) \\ &= \frac{1}{\mathbb{E}[B^{(n)}]}\int_0^\infty x^p\mathbb{P}(B^{(n)} \le x)\,\mathrm{d}x \\ &= \frac{1}{(p+1)\mathbb{E}[B^{(n)}]}\int_0^\infty x^{p+1}\,\mathrm{d}\mathbb{P}(B^{(n)} \le x) \\ &= \frac{1}{(p+1)\mathbb{E}[B^{(n)}]}\mathbb{E}[(B^{(n)})^{p+1}].\end{aligned} \tag{6.7}$$

Therefore, we need to show that $(\mu^{(n)})^p\mathbb{E}[(K_1^{(n)})^p]$ is uniformly bounded as $n \to \infty$. Let $k = \lfloor p \rfloor$. Then

$$\begin{aligned}\mathbb{E}[(K_1^{(n)})^p] &= \frac{1-\xi_1^{(n)}}{(1-\xi_1^{(n)})^p}\sum_{m=0}^\infty ((1-\xi_1^{(n)})m)^p(\xi_1^{(n)})^m \\ &\le \frac{1-\xi_1^{(n)}}{(1-\xi_1^{(n)})^p}\sum_{m=0}^{\left\lfloor \frac{1}{1-\xi_1^{(n)}}\right\rfloor} ((1-\xi_1^{(n)})m)^k(\xi_1^{(n)})^m \\ &\quad + \frac{1-\xi_1^{(n)}}{(1-\xi_1^{(n)})^p}\sum_{m=\left\lfloor \frac{1}{1-\xi_1^{(n)}}\right\rfloor+1}^\infty ((1-\xi_1^{(n)})m)^{k+1}(\xi_1^{(n)})^m \\ &\le \frac{(1-\xi_1^{(n)})^{k+1}}{(1-\xi_1^{(n)})^p}\sum_{m=0}^\infty m^k(\xi_1^{(n)})^m + \frac{(1-\xi_1^{(n)})^{k+2}}{(1-\xi_1^{(n)})^p}\sum_{m=0}^\infty m^{k+1}(\xi_1^{(n)})^m.\end{aligned} \tag{6.8}$$

On the one hand, for any $\ell \in \{1, 2, \ldots\}$, we have

$$(1-\xi_1^{(n)})^{\ell+1}(\xi_1^{(n)})^\ell \frac{\mathrm{d}^\ell}{\mathrm{d}(\xi_1^{(n)})^\ell} \sum_{m=0}^{\infty} (\xi_1^{(n)})^m = (1-\xi_1^{(n)})^{\ell+1}(\xi_1^{(n)})^\ell \frac{\mathrm{d}^\ell}{\mathrm{d}(\xi_1^{(n)})^\ell}(1-\xi_1^{(n)})^{-1}$$
$$= \ell!(\xi_1^{(n)})^\ell. \tag{6.9}$$

On the other hand we have

$$(1-\xi_1^{(n)})^{\ell+1}(\xi_1^{(n)})^\ell \frac{\mathrm{d}^\ell}{\mathrm{d}(\xi_1^{(n)})^\ell} \sum_{m=0}^{\infty} (\xi_1^{(n)})^m = (1-\xi_1^{(n)})^{\ell+1} \sum_{m=\ell}^{\infty} m(m-1)\cdots(m-\ell+1)(\xi_1^{(n)})^m$$
$$= (1-\xi_1^{(n)})^{\ell+1} \sum_{m=0}^{\infty} m^\ell (\xi_1^{(n)})^m$$
$$- (1-\xi_1^{(n)})^{\ell+1} \sum_{m=0}^{\ell-1} m^\ell (\xi_1^{(n)})^m$$
$$+ (1-\xi_1^{(n)})^{\ell+1} \sum_{m=\ell}^{\infty} o(m^\ell)(\xi_1^{(n)})^m. \tag{6.10}$$

Combining equalities (6.9) and (6.10), we find that

$$(1-\xi_1^{(n)})^{\ell+1} \sum_{m=0}^{\infty} m^\ell (\xi_1^{(n)})^m = \ell!(\xi_1^{(n)})^\ell + (1-\xi_1^{(n)})^{\ell+1} \sum_{m=0}^{\ell-1} m^\ell (\xi_1^{(n)})^m$$
$$+ (1-\xi_1^{(n)})^{\ell+1} \sum_{m=\ell}^{\infty} o(m^\ell)(\xi_1^{(n)})^m.$$

Now for any $\nu > 0$ there exists a $M_\nu \in \mathbb{N}$ such that for all $m \geq M_\nu$ the $o(m)$ term in the last expression is dominated by νm^ℓ. Fix such $\nu \in (0,1)$ and M_ν. Then

$$(1-\xi_1^{(n)})^{\ell+1} \sum_{m=0}^{\infty} m^\ell (\xi_1^{(n)})^m \leq \ell! + \ell^{\ell+1} + (1-\xi_1^{(n)})^{\ell+1} \nu \sum_{m=M_\nu}^{\infty} m^\ell (\xi_1^{(n)})^m$$
$$+ (1-\xi_1^{(n)})^{\ell+1} \sum_{m=0}^{M_\nu} o(m^\ell)(\xi_1^{(n)})^m$$
$$\leq C_6 + (1-\xi_1^{(n)})^{\ell+1} \nu \sum_{m=0}^{\infty} m^\ell (\xi_1^{(n)})^m$$

for some constant $C_6 > 0$, and hence

$$(1-\xi_1^{(n)})^{\ell+1} \sum_{m=0}^{\infty} m^\ell (\xi_1^{(n)})^m \leq \frac{C_6}{1-\nu}. \tag{6.11}$$

Since $(\mu^{(n)}/(1-\xi_1^{(n)}))^p = (\mathbb{E}[A^{(n)}] + \mathbb{E}[B^{(n)}])^p$, we may conclude from (6.8) and (6.11) that $(\mu^{(n)})^p \mathbb{E}[(K_1^{(n)})^p]$ is uniformly bounded from above as $n \to \infty$, and so is $(\mu^{(n)})^p \mathbb{E}[(W_1^{(n)})^p]$ by (6.6).

Second, we consider the functional $W_2^{(n)}$. Recall that $W_2^{(n)}$ denotes the steady-state workload in a GI/M/1 system upon arrival. Theorem VIII.5.8 and page 296 in Asmussen [2] together state that

$$W_2^{(n)} \stackrel{d}{=} \sum_{i=1}^{K_2^{(n)}} E_i^{(n)}, \tag{6.12}$$

where $K_2^{(n)}$ is a geometrically distributed random variable with support $\{0, 1, \ldots\}$ and unknown fail parameter $\xi_2^{(n)}$. Remerova et al. [25] again ensure

$$\mathbb{E}[(W_2^{(n)})^p] \leq \mathbb{E}[(K_2^{(n)})^p]\mathbb{E}[(E^{(n)})^p], \tag{6.13}$$

where the latter expectation is finite uniformly in n as a property of exponential distributions. The p-th moment of $K_2^{(n)}$ is bounded by (6.8) and (6.11), so that

$$(\mu^{(n)})^p\mathbb{E}[(K_2^{(n)})^p] \leq C_7 \left(\frac{\mu^{(n)}}{1-\xi_2^{(n)}}\right)^p \tag{6.14}$$

for some constant $C_7 > 0$. The proof is complete if we show $\frac{\mu^{(n)}}{1-\xi_2^{(n)}} = O(1)$.

From (6.12) it can be verified that $\mathbb{P}(W_2^{(n)} = 0) = \mathbb{P}(K_2^{(n)} = 0) = 1 - \xi_2^{(n)}$. Additionally, by Theorem VIII.2.3 in Asmussen [2], we have $\mathbb{P}(W_2^{(n)} = 0) = 1/\mathbb{E}[N_2^{(n)}]$, thereby implying $1 - \xi_2^{(n)} = 1/\mathbb{E}[N_2^{(n)}]$. Here, $N_2^{(n)}$ is the steady-state number of jobs in a busy period of the GI/M/1 queue. Applying Lemma 1 with interarrival times $A_i^{(n)}$, job sizes $E_i^{(n)}$ and mean service between arrivals $\frac{1}{2}\mu^{(n)}$ yields $\frac{1}{2}\mu^{(n)}\mathbb{E}[N_2^{(n)}] = \Theta(1)$, and therefore in particular $\frac{\mu^{(n)}}{1-\xi_2^{(n)}} = O(1)$.

Finally, for $0 < p < 1$ the lemma follows from the case $p = 1$ after observing that $(\mu^{(n)})^p\mathbb{E}[(W^{(n)})^p] \leq \left(\mu^{(n)}\mathbb{E}[W^{(n)}]\right)^p$ by Jensen's inequality.

□

Lemmas 1 and 2 provide the asymptotic behaviour of some functionals that are closely related to $P^{(n)}$ and $N^{(n)}$. The remainder of this section utilizes these results in order to prove Propositions 1 and 2.

6.1 Busy period duration $P^{(n)}$

This section is devoted to the proof of Proposition 1. We wish to show that

$$\mathbb{E}[(P^{(n)})^\kappa] = O\left((1-\rho^{(n)})^{-(2\kappa-1)}\right). \tag{5.7, revisited}$$

for all $\kappa \in [1, \alpha]$, provided that $\sup_{n=1,2,\ldots} \mathbb{E}[(B^{(n)})^\alpha] < \infty$ for some $\alpha \geq 2$. Moreover, $\mathbb{E}[P^{(n)}] = \Theta\left((1-\rho^{(n)})^{-1}\right)$.

First, consider $\kappa = 1$. Due to Little's law for a busy server, we have

$$1 - \rho^{(n)} = \frac{\mathbb{E}[I^{(n)}]}{\mathbb{E}[I^{(n)}] + \mathbb{E}[P^{(n)}]}, \tag{6.15}$$

so that

$$\mathbb{E}[P^{(n)}] = \frac{\rho^{(n)}\mathbb{E}[I^{(n)}]}{1-\rho^{(n)}}. \tag{6.16}$$

The result now follows from Lemma 1.

Second, consider $\kappa > 1$. Using renewal theory, we have [2, Theorem V.4.6]

$$\mathbb{E}[(P_e^{(n)})^{\kappa-1}] = \frac{\mathbb{E}[(P^{(n)})^\kappa]}{\kappa\mathbb{E}[P^{(n)}]}. \tag{6.17}$$

Thus, it suffices to show that $\mathbb{E}[(P_e^{(n)})^{\kappa-1}] = O((1-\rho^{(n)})^{-2(\kappa-1)})$. We have the following convenient representation for $P_e^{(n)}$ [2, Theorem X.3.4]:

$$\begin{aligned} P_e^{(n)} &\stackrel{d}{=} \inf\{\tau \geq 0 : \widetilde{X}^{(n)}(\tau) \leq -V^{(n)} \mid V^{(n)} > 0\} \\ &\stackrel{d}{=} \inf\{\tau \geq 0 : B_e^{(n)} + W^{(n)} + \widetilde{X}^{(n)}(\tau) \leq 0\}. \end{aligned}$$

The above relation allows us to bound $\mathbb{P}(P_e^{(n)} > t)$ as

$$\begin{aligned}
\mathbb{P}(P_e^{(n)} > t) &= \mathbb{P}(\inf\{\tau \geq 0 : B_e^{(n)} + W^{(n)} + \widetilde{X}^{(n)}(\tau) \leq 0\} > t) \\
&= \mathbb{P}(B_e^{(n)} + W^{(n)} + \widetilde{X}^{(n)}(\tau) > 0, \forall \tau \leq t) \\
&\leq \mathbb{P}(B_e^{(n)} + W^{(n)} + \widetilde{X}^{(n)}(t) + (1-\rho^{(n)})t/2 > (1-\rho^{(n)})t/2) \\
&\leq \mathbb{P}(B_e^{(n)} > (1-\rho^{(n)})t/6) + \mathbb{P}(W^{(n)} > (1-\rho^{(n)})t/6) \\
&\qquad + \mathbb{P}\left(\sup_{\tau\geq 0}[\widetilde{X}^{(n)}(\tau) + (1-\rho^{(n)})\tau/2] > (1-\rho^{(n)})t/6\right).
\end{aligned}$$

In Section 2 we derived the relations $\widetilde{X}^{(n)}(\tau) \leq_{st} X^{(n)}(\tau)$ and $W^{(n)} \stackrel{d}{=} \sup_{\tau\geq 0} X^{(n)}(\tau)$. These relations imply

$$\begin{aligned}
\mathbb{P}(P_e^{(n)} > t) \leq \mathbb{P}(B_e^{(n)} > (1-\rho^{(n)})t/6) \\
+ 2\mathbb{P}\left(\sup_{\tau\geq 0}[X^{(n)}(\tau) + (1-\rho^{(n)})\tau/2] > (1-\rho^{(n)})t/6\right). \quad (6.18)
\end{aligned}$$

The final inequality now suggests that the expectation $\mathbb{E}[(P_e^{(n)})^{\kappa-1}]$ is dominated by

$$\begin{aligned}
\mathbb{E}[(P_e^{(n)})^{\kappa-1}] &= (\kappa-1)\int_0^\infty t^{\kappa-2}\mathbb{P}(P_e^{(n)} > t)\,\mathrm{d}t \\
&\leq (\kappa-1)\int_0^\infty t^{\kappa-2}\mathbb{P}(B_e^{(n)} > (1-\rho^{(n)})t/6)\,\mathrm{d}t \\
&\qquad + 2(\kappa-1)\int_0^\infty t^{\kappa-2}\mathbb{P}\left(\sup_{\tau\geq 0}[X^{(n)}(\tau) + (1-\rho^{(n)})\tau/2] > (1-\rho^{(n)})t/6\right)\,\mathrm{d}t.
\end{aligned}$$

To deal with the first term, note that

$$\begin{aligned}
(\kappa-1)\int_0^\infty t^{\kappa-2}\mathbb{P}(B_e^{(n)} > (1-\rho^{(n)})t/6)\,\mathrm{d}t &= (\kappa-1)\int_0^\infty t^{\kappa-2}\mathbb{P}(6B_e^{(n)}/(1-\rho^{(n)}) > t)\,\mathrm{d}t \\
&= \mathbb{E}\left[\left(\frac{6B_e^{(n)}}{1-\rho^{(n)}}\right)^{\kappa-1}\right] \\
&= O\left((1-\rho^{(n)})^{-(\kappa-1)}\right)\mathbb{E}[(B_e^{(n)})^{\kappa-1}] \\
&= o\left((1-\rho^{(n)})^{-2(\kappa-1)}\right),
\end{aligned}$$

since $\mathbb{E}[(B^{(n)})^\kappa] < \infty$ implies $\mathbb{E}[(B_e^{(n)})^{\kappa-1}] < \infty$ (cf. 6.7).

For the second term, observe that

$$\begin{aligned}
\sup_{\tau\geq 0}[X^{(n)}(\tau) + (1-\rho^{(n)})\tau/2] &= \sup_{\tau\geq 0}\left[\sum_{i=1}^{N^{(n)}(\tau)} B_i^{(n)} - \tau + (1-\rho^{(n)})\tau/2\right] \\
&= \sup_{\tau\geq 0}\left[\sum_{i=1}^{N^{(n)}(\tau)} B_i^{(n)} - \frac{1+\rho^{(n)}}{2}\tau\right] \\
&= \sup_{n=0,1,2,\ldots}\left[\sum_{i=1}^{n}\left\{B_i^{(n)} - \frac{1+\rho^{(n)}}{2}A_i^{(n)}\right\}\right] \\
&=: \widetilde{W}^{(n)},
\end{aligned}$$

where $\widetilde{W}^{(n)}$ is equal in distribution to the steady-state cumulative amount of work at arrival in a GI/GI/1 queue with job sizes $B_i^{(n)}$ and interarrival times $\frac{1+\rho^{(n)}}{2}A_i^{(n)}, i \geq 1$. Define $\widetilde{\mu}^{(n)} := \frac{1+\rho^{(n)}}{2}\mathbb{E}[A^{(n)}] - \mathbb{E}[B^{(n)}] = \frac{1-\rho^{(n)}}{2}\mathbb{E}[A^{(n)}]$. The second integral may yet be rewritten as

$$\begin{aligned}(\kappa-1)\int_0^\infty t^{\kappa-2}\mathbb{P}\left(\sup_{\tau\geq 0}[X^{(n)}(\tau)+(1-\rho^{(n)})\tau/2] > (1-\rho^{(n)})t/6\right)\,\mathrm{d}t \\ &= (\kappa-1)\int_0^\infty t^{\kappa-2}\mathbb{P}\left(\frac{6}{1-\rho^{(n)}}\widetilde{W}^{(n)} > t\right)\,\mathrm{d}t \\ &= \frac{6^{\kappa-1}}{(1-\rho^{(n)})^{\kappa-1}(\widetilde{\mu}^{(n)})^{\kappa-1}}\mathbb{E}[(\widetilde{\mu}^{(n)}\widetilde{W}^{(n)})^{\kappa-1}] \\ &\leq \frac{C_8}{(1-\rho^{(n)})^{2(\kappa-1)}}\mathbb{E}[(\widetilde{\mu}^{(n)}\widetilde{W}^{(n)})^{\kappa-1}]\end{aligned}$$

for some constant $C_8 > 0$ by Lemma 1. All that is left to establish is that $\mathbb{E}[(\widetilde{\mu}^{(n)}\widetilde{W}^{(n)})^{\kappa-1}]$ is bounded, but this follows from Lemma 2.

6.2 Arrivals in a busy period $N^{(n)}$

This section contains the proof of Proposition 2. The proposition states that

$$\mathbb{E}[(N^{(n)})^\kappa] = O\left((1-\rho^{(n)})^{-(2\kappa-1)}\right). \qquad \text{(5.8, revisited)}$$

for all $\kappa \in [1,\alpha]$, provided that $\sup_{n=1,2,\ldots}\mathbb{E}[(B^{(n)})^\alpha] < \infty$ for some $\alpha \geq 2$. Moreover, $\mathbb{E}[N^{(n)}] = \Theta\left((1-\rho^{(n)})^{-1}\right)$.

The structure of the proof is identical to the proof of Proposition 1. For $\kappa = 1$, the result follows directly from Lemma 1. Therefore, we consider $\mathbb{E}[(N^{(n)})^\kappa]$ for $\kappa > 1$. Similar to the proof of Proposition 1, we use the relation

$$\mathbb{E}[(N_e^{(n)})^{\kappa-1}] = \frac{\mathbb{E}[(N^{(n)})^\kappa]}{\kappa\mathbb{E}[N^{(n)}]} \qquad (6.19)$$

and note that

$$\begin{aligned}N_e^{(n)} &\stackrel{d}{=} \inf\{\eta\in\{0,1,2,\ldots\} : \widetilde{S}_\eta^{(n)} \leq -V^{(n)} \mid V^{(n)} > 0\} \\ &\stackrel{d}{=} \inf\{\eta\in\{0,1,2,\ldots\} : B_e^{(n)} + W^{(n)} + \widetilde{S}_\eta^{(n)} \leq 0\}.\end{aligned}$$

As before, the relations $\widetilde{S}_\eta^{(n)} \leq_{st} S_\eta^{(n)}$ and $W^{(n)} \stackrel{d}{=} \sup_{\eta=0,1,\ldots} S_\eta^{(n)}$ are exploited in order to obtain an equivalent of (6.18):

$$\begin{aligned}\mathbb{P}(N_e^{(n)} > m) \leq \mathbb{P}(B_e^{(n)} > (1-\rho^{(n)})m/6) \\ + 2\mathbb{P}\left(\sup_{\eta\in\{0,1,2,\ldots\}}[S_\eta^{(n)} + (1-\rho^{(n)})\eta/2] > (1-\rho^{(n)})m/6\right).\end{aligned}$$

For any $\eta \in \{0,1,2,\ldots\}$, consider $\tau_\eta^{(n)} := A_1^{(n)} + \ldots + A_\eta^{(n)}$. Then $S_\eta^{(n)} = X^{(n)}(\tau_\eta^{(n)})$, so in particular

$$\begin{aligned}\mathbb{P}(N_e^{(n)} > m) \leq \mathbb{P}(B_e^{(n)} > (1-\rho^{(n)})m/6) \\ + 2\mathbb{P}\left(\sup_{\tau\geq 0}[X^{(n)}(\tau) + (1-\rho^{(n)})\tau/2] > (1-\rho^{(n)})m/6\right).\end{aligned}$$

The remainder of the proof is identical to that of Proposition 1.

7 Conclusion

In this paper, we proved a result about the average case performance of (an extension of) the Randomized Multilevel Feedback (RMLF) algorithm in a GI/GI/1 queue. Specifically, the gap in average sojourn time between the RMLF algorithm and the Shortest Remaining Processing Time algorithm behaves like $O(\log(1/(1-\rho^{(n)})))$ and this bound is tight for the M/M/1 queue. An appealing property of the RMLF algorithm is that its implementation does not depend on the input distributions G_a and G_b; however, if G_a and G_b are known then there can be blind algorithms with a better performance than RMLF (e.g. Foreground-Background if G_b has decreasing failure rate). The result was established by using techniques from both competitive analysis and applied probability. As the structure of the proof is quite general, it would be interesting to explore other possibilities in the intersection of these areas.

Acknowledgements

We would like to thank Adam Wierman for his encouragements to work on this research topic.

A The eRMLF algorithm: definition and competitive ratio

The eRMLF algorithm is presented after the introduction of some notation. Define the queues $\widetilde{Q}_z, z \in \mathbb{Z}$ and a "new job" queue $\widetilde{Q}^*$. Let the targets $\widetilde{T}_{z,j}$ be given by $\widetilde{T}_{z,j} = 2^z \max\{1, 2-\widetilde{\beta}_j\}$, where $\widetilde{\beta}_j$ is an exponentially distributed random variable with $\mathbb{P}(\widetilde{\beta}_j \leq x) = 1 - \exp[-\theta x \ln j]$. Identical to the RMLF algorithm, θ is a symbolic constant fixed at $\theta = 12$. All symbols $\widetilde{J}_j, \widetilde{r}_j, \widetilde{B}_j$ and $\widetilde{w}_j(t)$ are defined analogue to the symbols without accent in the RMLF algorithm. All release times $\widetilde{r}_j$ must be distinct (e.g. all interarrival times are strictly positive), and jobs may have any size $\widetilde{B}_j \geq 0$.

Every job $\widetilde{J}_h$ is assigned a target $\widetilde{T}_{*,h}$ upon arrival, after which it is immediately served in $\widetilde{Q}^*$ by a dedicated server. It departs from $\widetilde{Q}^*$ on three occasions:

- The amount of service received equals the size $\widetilde{B}_h$ of the job. In this case, $\widetilde{J}_h$ is completed and leaves the system.
- A new job enters the system. In this case, $\widetilde{J}_h$ is moved to a queue $\widetilde{Q}_z, z \in \mathbb{Z}$, that it naturally belongs to based on the amount of service $\widetilde{w}_h(t)$ it has obtained thus far; that is, it is moved to the unique queue $\widetilde{Q}_{z^*}$ that satisfies $\widetilde{T}_{z^*-1,h} \leq \widetilde{w}_h(t) < \widetilde{T}_{z^*,h}$.
- The amount of service received equals the initial target $\widetilde{T}_{*,h}$. In this case, similar to the previous case, $\widetilde{J}_h$ is moved to a queue $\widetilde{Q}_z$ that it naturally belongs to.

The choice of the initial target $\widetilde{T}_{*,h}$ depends on the system state:

- If the system is empty upon arrival, then the server is dedicated to $\widetilde{J}_h$ regardless of the queue that $\widetilde{J}_h$ is in. In this case, the target can be chosen arbitrarily; we set it to $\widetilde{T}_{*,h} = \widetilde{T}_{0,h}$.
- If the system is not empty upon arrival, then there must be a lowest-index non-empty queue $\widetilde{Q}_{z^*}$ (possibly after moving the job originally in $\widetilde{Q}^*$ to another queue). $\widetilde{J}_h$ may now experience a dedicated server until the moment when it would enter queue $\widetilde{Q}_{z^*}$ based on its obtained service and the (z^*-1)-th target $\widetilde{T}_{z^*-1,h}$. Therefore, $\widetilde{J}_h$ should be moved no later then after $\widetilde{T}_{*,h} = \widetilde{T}_{z^*-1,h}$ units of obtained service.

If $\widetilde{Q}^*$ is empty, then eRMLF always works on the non-empty queue $\widetilde{Q}_z$ with the lowest index $z \in \mathbb{Z}$. Jobs in $\widetilde{Q}_z$ are always served in a First Come First Serve order.

The eRMLF is formally presented in Figure 2. Observe that both RMLF and eRMLF preserve the ordering of the jobs; that is, if job $\widetilde{J}_j$ is released prior to job $\widetilde{J}_k$ then as long as both jobs are incomplete:

- job $\widetilde{J}_j$ will never be in a lower queue than job $\widetilde{J}_k$, and
- if both jobs are in the same queue, then job $\widetilde{J}_j$ is closer to the front of the queue than job $\widetilde{J}_k$.

We are now ready to prove Theorem 2, stating that

$$\mathbb{E}[T_{\text{eRMLF}}(\mathcal{I})] \leq C_1 \log(m) T_{\text{SRPT}}(\mathcal{I}) \tag{3.3, revisited}$$

for all instances $\mathcal{I}$ of size at most m for a universal constant C_1. This constant is identical to the constant C_1 in Theorem 1.

Proof. Consider any instance $\widetilde{\mathcal{I}}$ for eRMLF of size at most m. Since all jobs of size zero are immediately served in queue $\widetilde{Q}^*$ upon arrival, we assume without loss of generality that the instance does not contain any jobs of size zero. Consequently, the minimum job size $\widetilde{B}_{\min} = \min_{j=1,\ldots,|\widetilde{\mathcal{I}}|} \widetilde{B}_j$ is strictly positive. We now transform the instance $\widetilde{\mathcal{I}}$ for eRMLF to a corresponding instance $\mathcal{I}$ for RMLF.

Define the scaling parameter $g := \lfloor \log_2(\widetilde{B}_{\min}) \rfloor - 1 \in \mathbb{Z}$, satisfying $2^{-g}\widetilde{B}_{\min} \geq 2$. The instance $\mathcal{I}$ consists out of $|\widetilde{\mathcal{I}}|$ jobs that are scaled versions of the original $|\widetilde{\mathcal{I}}|$ jobs; specifically, job J_j has size $B_j := 2^{-g}\widetilde{B}_j$ and release date $r_j := 2^{-g}\widetilde{r}_j$. In particular, the smallest job is of size at least 2 and the RMLF algorithm may be applied to the instance $\mathcal{I}$.

Since the jobs are released in the same order as in the original instance, we note that the random variables β_j assigned by RMLF have the same distribution as the $\widetilde{\beta}_j$ assigned by eRMLF. We therefore couple these random variables in a trivial way: $\beta_j \equiv \widetilde{\beta}_j$ for all $j = 1, \ldots, |\widetilde{\mathcal{I}}|$. It immediately follows that the targets $\widetilde{T}_{z,j}$ as assigned to $\widetilde{\mathcal{I}}$ by eRMLF and the targets $T_{i,j}$ as assigned to $\mathcal{I}$ by RMLF satisfy $\widetilde{T}_{i,j} = T_{i,j}$ for all $i \geq 0$. Additionally, we have $T_{0,j} = \widetilde{T}_{0,j} = 2^{g-g} \max\{1, 2 - \widetilde{\beta}_j\} = 2^{-g}\widetilde{T}_{g,j}$.

We will show that the above construction implies an equivalence between RMLF and eRMLF. For all $z \in \mathbb{Z}$ and $t \geq 0$ define the sets

$$\widetilde{Q}_z(t) := \left\{ \widetilde{J}_j : \widetilde{T}_{z-1,j} \leq \widetilde{w}_j(t) < \widetilde{T}_{z,j} \right\} \tag{A.1}$$

that contain all jobs in the system at time t that are in queue $\widetilde{Q}_z$, *or the most recently released job if it has received a similar amount of service.* The equivalence is first observed between the initial RMLF queue $Q_0(t)$ and the augmented eRMLF queue $\widehat{Q}(t)$, defined as

$$Q_0(t) := \{ J_j : w_j(t) < T_{0,j} \} \tag{A.2}$$

and

$$\widehat{Q}(t) := \bigcup_{z=-\infty}^{g} \widetilde{Q}_z(t), \tag{A.3}$$

respectively. One may observe that since Q_0 is the highest priority queue, it experiences a dedicated, work-conserving server that works at unit speed on jobs J_j with sizes $T_{0,j}$. Therefore, the event that the RMLF server works on Q_0 is equivalent to the event $\{Q_0(t) > 0\}$.

We assume without loss of generality that job 1 arrives at time $r_1 = 0$. The arrival of this job initiates a first busy period for $Q_0(t)$ of N_1 jobs, where N_1 is such that the cumulative

Algorithm eRMLF**:** At all times the collection of released, but uncompleted, jobs are partitioned into queues, $\widetilde{Q}^*, \widetilde{Q}_z, z \in \mathbb{Z}$. We say that $\widetilde{Q}_i$ is lower than $\widetilde{Q}_j$ for $i < j$. $\widetilde{Q}^*$ is the lowest queue. For each job $\widetilde{J}_j \in \widetilde{Q}_i, \widetilde{T}_{i,j} \in [2^i, 2^{i+1}]$ when it entered $\widetilde{Q}_i$. eRMLF maintains the invariant that it is always running the job at the front of the lowest non-empty queue.
When a job $\widetilde{J}_h$ is released at time $\widetilde{r}_h$, eRMLF takes the following actions:

- If, just prior to $\widetilde{r}_h$, all queues were empty, then
  - Job $\widetilde{J}_h$ is enqueued on $\widetilde{Q}^*$.
  - The target $\widetilde{T}_{*,h}$ is set to $\widetilde{T}_{0,h} = \max\{1, 2 - \widetilde{\beta}_h\}$.
- If, just prior to $\widetilde{r}_h$, there are unfinished jobs in the system but $\widetilde{Q}^*$ is empty, then
  - Job $\widetilde{J}_h$ is enqueued on $\widetilde{Q}^*$.
  - The target $\widetilde{T}_{*,h}$ is set to $\widetilde{T}_{z^*-1,h} = 2^{z^*-1}\max\{1, 2 - \widetilde{\beta}_h\}$, where $z^* = \min\{z \in \mathbb{Z} : \widetilde{Q}_z \text{ non-empty at time } t\}$ is the index of the lowest non-empty queue.
- If, just prior to $\widetilde{r}_h$, $\widetilde{Q}^*$ is non-empty, then $\widetilde{Q}^* = \{\widetilde{J}_{h-1}\}$ at that time. Now,
  - The target $\widetilde{T}_{z^*,h-1} = 2^{z^*}\max\{1, 2 - \widetilde{\beta}_{h-1}\}$ with $z^* := \min\{z \in \mathbb{Z} : \widetilde{w}_{h-1}(\widetilde{r}_h) \le \widetilde{T}_{z,h-1}\}\}$ is the lowest target not yet reached by job $\widetilde{J}_{h-1}$.
  - Job $\widetilde{J}_{h-1}$ is dequeued from $\widetilde{Q}^*$.
  - Job $\widetilde{J}_{h-1}$ is enqueued on $\widetilde{Q}_{z^*}$.
  - Job $\widetilde{J}_h$ is enqueued on $\widetilde{Q}^*$.
  - The target $\widetilde{T}_{*,h}$ is set to $\widetilde{T}_{z^*-1,h} = 2^{z^*-1}\max\{1, 2 - \widetilde{\beta}_h\}$.
- If, just prior to $\widetilde{r}_h$, it was the case that eRMLF was running a job $\widetilde{J}_j$, then $\widetilde{J}_j$ is pre-empted.
- eRMLF begins running $\widetilde{J}_h$.

If at some time t, a job $\widetilde{J}_j \in \widetilde{Q}_{z-1}$ is being run when $\widetilde{w}_j(t)$ becomes equal to $\widetilde{T}_{z-1,j}$, then eRMLF takes the following actions:

- Job $\widetilde{J}_j$ is dequeued from $\widetilde{Q}_{z-1}$.
- Job $\widetilde{J}_j$ is enqueued on $\widetilde{Q}_z$.
- The target $\widetilde{T}_{z,j}$ is set to $2\widetilde{T}_{z-1,j} = 2^i\max\{1, 2 - \widetilde{\beta}_j\}$.

If at some time t, a job $\widetilde{J}_j \in \widetilde{Q}^*$ is being run when $\widetilde{w}_j(t)$ becomes equal to $\widetilde{T}_{*,j}$, then eRMLF takes the following actions:

- Job $\widetilde{J}_j$ is dequeued from $\widetilde{Q}^*$.
- Job $\widetilde{J}_j$ is enqueued on $\widetilde{Q}_{z^*}$, where $z^* = \log_2\left(\widetilde{w}_j(t)/\max\{1, 2 - \widetilde{\beta}_j\}\right) + 1$.
- The target $\widetilde{T}_{z^*,j}$ is set to $2\widetilde{T}_{*,j} = 2^{z^*}\max\{1, 2 - \widetilde{\beta}_j\}$.

Whenever a job is completed, it is removed from its queue. Recall that eRMLF maintains the invariant that it is always running the job at the front of the lowest non-empty queue.

Figure 2: Formal statement of eRMLF algorithm.

sizes of the first N_1 jobs can be served before the (N_1+1)-th job is released. It is defined as $N_1 = \inf\{k \geq 1 : \sum_{j=1}^{k} T_{0,j} - r_{j+1} \leq 0\}$, where $r_{|\widetilde{\mathcal{I}}|+1}$ is understood as plus infinity. The duration of the busy period is given by $P_1 = \sum_{j=1}^{N_1} T_{0,j} = \sum_{j=1}^{N_1} 2^{-g} \widetilde{T}_{g,j}$. The server may then work on jobs in higher queues (perceived as idle time by Q_0) until time r_{N_1+1}, when a new busy period is initiated. For $t \in [0, r_{N_1+1})$ we have now obtained

$$Q_0(t) > 0 \Leftrightarrow t \leq P_1 \Leftrightarrow 2^g t \leq \sum_{j=1}^{N_1} \widetilde{T}_{g,j}. \tag{A.4}$$

By a similar analysis of the augmented queue $\widehat{Q}$ we then find $\widehat{Q}(t) > 0 \Leftrightarrow Q_0(2^{-g}t) > 0$ for $t \in [0, 2^g r_{N_1+1}) = [0, \widetilde{r}_{N_1+1})$, and for all $t \geq 0$ by a straightforward generalisation of the above procedure. Observing that both algorithms preserve the ordering of jobs, we may similarly show that the eRMLF server processes job $\widetilde{J}_j$ in queue $\widetilde{Q}_{g+i}$ at time t if and only if the RMLF server processes job J_j in queue Q_i at time $2^g t$ for all $i = 1, 2, \ldots$ and $t \geq 0$.

From the above results, one may deduce that the average sojourn time $\mathbb{E}[T_{\mathrm{eRMLF}}(\widetilde{\mathcal{I}})]$ of instance $\widetilde{\mathcal{I}}$ under algorithm eRMLF equals 2^g times the average sojourn time $\mathbb{E}[T_{\mathrm{RMLF}}(\mathcal{I})]$ of instance $\mathcal{I}$ under RMLF. The competitive ratio of RMLF as stated in Theorem 1 hence guarantees that, for all instances $\widetilde{\mathcal{I}}$ of size at most m,

$$\mathbb{E}[T_{\mathrm{eRMLF}}(\widetilde{\mathcal{I}})] = 2^g \mathbb{E}[T_{\mathrm{RMLF}}(\mathcal{I})] \leq C_1 \log(m) 2^g \mathbb{E}[T_{\mathrm{SRPT}}(\mathcal{I})]. \tag{A.5}$$

The competitive ratio of eRMLF is concluded by verifying

$$2^g \mathbb{E}[T_{\mathrm{SRPT}}(\mathcal{I})] = \mathbb{E}[T_{\mathrm{SRPT}}(\widetilde{\mathcal{I}})], \tag{A.6}$$

which is a direct consequence of our scaling. In particular, the constant C_1 in the upper bound is the same for RMLF and eRMLF. □

References

[1] Aalto, S. and Ayesta, U. (2006). Mean delay analysis of multi level processor sharing disciplines. In *INFOCOM*, volume 2006, pages 1–11.

[2] Asmussen, S. (2003). *Applied probability and queues*, volume 51. Springer.

[3] Bansal, N. (2005). On the average sojourn time under M/M/1/SRPT. *Operations Research Letters*, 33(2):195–200.

[4] Bansal, N. and Gamarnik, D. (2006). Handling load with less stress. *Queueing Systems*, 54(1):45–54.

[5] Bansal, N. and Harchol-Balter, M. (2001). Analysis of SRPT scheduling: Investigating unfairness. In *Proceedings of the 2001 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*, SIGMETRICS '01, pages 279–290, New York, NY, USA. ACM.

[6] Becchetti, L. and Leonardi, S. (2001). Non-clairvoyant scheduling to minimize the average flow time on single and parallel machines. In *Proceedings of the Thirty-third Annual ACM Symposium on Theory of Computing*, STOC '01, pages 94–103, New York, NY, USA. ACM.

[7] Becchetti, L. and Leonardi, S. (2004). Nonclairvoyant scheduling to minimize the total flow time on single and parallel machines. *J. ACM*, 51(4):517–539.

[8] Borodin, A. and El-Yaniv, R. (1998). *Online computation and competitive analysis*. Cambridge University Press.

[9] Boxma, O. and Zwart, B. (2007). Tails in scheduling. *SIGMETRICS Perform. Eval. Rev.*, 34(4):13–20.

[10] Conway, R. W., Maxwell, W. L., and Miller, L. W. (1967). *Theory of scheduling*. Addison-Wesley Publishing Company, Inc.

[11] Corbató, F. J., Merwin-Daggett, M., and Daley, R. C. (1962). An experimental time-sharing system. In *Proceedings of the May 1-3, 1962, Spring Joint Computer Conference*, AIEE-IRE '62 (Spring), pages 335–344, New York, NY, USA. ACM.

[12] Fiat, A. and Woeginger, G. J. (1998). Online algorithms: The state of the art. *LNCS, Springer-Verlag*, 1442.

[13] Gromoll, H. C. (2004). Diffusion approximation for a processor sharing queue in heavy traffic. *Annals of Applied Probability*, 14(2):555–611.

[14] Kalyanasundaram, B. and Pruhs, K. R. (2003). Minimizing flow time nonclairvoyantly. *Journal of the ACM (JACM)*, 50(4):551–567.

[15] Kleinrock, L. (1976). *Queueing Systems II: Computer Theory*. John Wiley & Sons.

[16] Lin, M., Wierman, A., and Zwart, B. (2011). Heavy-traffic analysis of mean response time under shortest remaining processing time. *Performance Evaluation*, 68(10):955–966.

[17] Lotov, V. I. (2002). Inequalities for the moments and distribution of the ladder height of a random walk. *Siberian Mathematical Journal*, 43(4):655–660.

[18] Lu, D., Sheng, H., and Dinda, P. (2004). Size-based scheduling policies with inaccurate scheduling information. In *Modeling, Analysis, and Simulation of Computer and Telecommunications Systems, 2004. (MASCOTS 2004). Proceedings. The IEEE Computer Society's 12th Annual International Symposium on*, pages 31–38.

[19] Motwani, R., Phillips, S., and Torng, E. (1994). Nonclairvoyant scheduling. *Theoretical Computer Science*, 130(1):17–47.

[20] Nair, J., Wierman, A., and Zwart, B. (2010). Tail-robust scheduling via limited processor sharing. *Performance Evaluation*, 67(11):978–995. Performance 2010.

[21] Nuyens, M. and Wierman, A. (2008). The foreground–background queue: a survey. *Performance evaluation*, 65(3):286–307.

[22] Nuyens, M., Wierman, A., and Zwart, B. (2008). Preventing large sojourn times using SMART scheduling. *Operations Research*, 56(1):88–101.

[23] Pruhs, K., Sgall, J., and Torng, E. (2004). Online scheduling. *Handbook of scheduling: algorithms, models, and performance analysis*.

[24] Puha, A. L. (2014). Diffusion limits for shortest remaining processing time queues under nonstandard spatial scaling. `http://arxiv.org/abs/1407.3837`.

[25] Remerova, M., Foss, S., and Zwart, B. (2014). Random fluid limit of an overloaded polling model. *Advances in Applied Probability*, 46(1):76–101.

[26] Righter, R., Shanthikumar, J. G., and Yamazaki, G. (1990). On extremal service disciplines in single-stage queueing systems. *Journal of Applied Probability*, 27(2):409–416.

[27] Schrage, L. (1968). Letter to the editor — a proof of the optimality of the shortest remaining processing time discipline. *Operations Research*, 16(3):687–690.

[28] Stolyar, A. L. and Ramanan, K. (2001). Largest weighted delay first scheduling: Large deviations and optimality. *The Annals of Applied Probability*, 11(1):1–48.

[29] Wierman, A., Harchol-Balter, M., and Osogami, T. (2005). Nearly insensitive bounds on SMART scheduling. In *Proceedings of the 2005 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*, SIGMETRICS '05, pages 205–216, New York, NY, USA. ACM.

[30] Wierman, A. and Nuyens, M. (2008). Scheduling despite inexact job-size information. *SIGMETRICS Perform. Eval. Rev.*, 36(1):25–36.

[31] Wierman, A. and Zwart, A. P. (2012). Is tail-optimal scheduling possible? *Operations Research*, 60(5):1249–1257.